

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Welch

Application No.: 10/785,434

Art Unit.: 4121

Filed: 2/24/2004

Examiner: Keehn, Richard

For: DISTRIBUTED MONITORING IN A TELECOMMUNICATIONS SYSTEM

Mail Stop Appeal Brief - Patents

Commissioner for Patents

P. O. Box 1450

Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

The Appellants herewith file this Appeal Brief in support of their appeal in the above identified matter.

TABLE OF CONTENTS

Item	Page Numbers
Identification page	1
Table of contents	2
(i) Real party in interest	3
(ii) Related appeals and interferences	3
(iii) Status of claims	3
(iv) Status of amendments	3
(v) Summary of claimed subject matter	4
(vi) Grounds of rejection to be reviewed on appeal	6
(vii) Argument	7
(viii) Claims appendix	14
(ix) Evidence appendix	18
(x) Related proceedings appendix	18
Summary	19

i. REAL PARTY IN INTEREST

The real party in interest is Lucent Technologies Inc., which is the assignee of the patent rights in the above-identified matter. The inventor, Arthur Welch, assigned the patent rights to Lucent Technologies Inc., which is recorded at reel/frame 015021/0501.

ii. RELATED APPEALS AND INTERFERENCES

No other appeals, interferences, or related applications are known to the Appellants, the Appellants' legal representative, or the Assignee, which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

iii. STATUS OF CLAIMS

Claims 1, 5-6, 8-11, 15-16 and 18-20 stand rejected and remain in the application for consideration on appeal. Claims 2-4, 7, 12-14 and 17 are cancelled.

iv. STATUS OF AMENDMENTS

No amendments have been filed since the response to the Office action mailed April 24, 2009.

v. **SUMMARY OF THE CLAIMED SUBJECT MATTER**

The following summarizes the independent claims presently pending:

Independent claim 1 recites a telecommunication system (100) configured to provide distributed system monitoring. See generally, FIG. 1. The telecommunication system includes a control system (110) and a plurality of peer communication devices (101-105). See FIG. 1. Responsive to handling telecommunications data, each peer communication device collects performance data, and transfers the performance data to the control system. See FIG. 2A and pages 4-5, lines 30-4. Responsive to receipt of the performance data from the peer communication devices, the control system processes the performance data from each of the peer communication devices to generate a performance file that indicates the performance of each of the peer communication devices, and transfers the performance file to each of the peer communication devices. See FIG. 2A and page 5, lines 5-15. Responsive to receipt of the performance file, each of the peer communication devices processes the performance file to compare its performance to the performance of the other peer communication devices to detect a fault. See FIG. 2A and page 5, lines 16-22. Responsive to detection of the fault, at least one of the peer communication devices processes the performance file to identify at least one recovery action, and performs the at least one recovery action to attempt to cure the fault. See FIG. 2B and pages 5-6, lines 23-6.

Independent claim 11 recites a method of operating a telecommunication system to provide distributed system monitoring, where the telecommunication system comprises a plurality of peer communication devices (101-105) coupled to a control system (110). See FIGS. 1 and 2A. The method includes collecting (step 202) performance data in each of the peer communication devices responsive to the peer communication devices handling telecommunications data, and transferring the performance data from each of the peer communication devices to the control system. See FIG. 2A and pages 4-5, lines 30-4. The method further includes processing, in the control system, (step 204) the performance data from each of the peer communication devices to generate a performance file that indicates the performance of each of the peer communication devices, and transferring the performance file from the control system to each of the peer communication devices. See FIG. 2A and page 5,

lines 5-15. The method further includes processing (step 206) the performance file in each of the peer communication devices to compare its performance to the performance of the other peer communication devices to detect a fault. See FIG. 2A and page 5, lines 16-22. Responsive to detection of the fault, the method further includes processing the performance file to identify at least one recovery action, and performing the at least one recovery action to attempt to cure the fault. See FIG. 2B and pages 5-6, lines 23-6.

vi. GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The Examiner rejected claims 1, 8-11, and 18-20 under 35 USC § 103(a) as being obvious in view of U.S. Patent Application Publication 2005/0065753 (Bigus) and U.S. Patent Application Publication 2003/0039352 (El-Fakih). The Examiner also rejected claims 5, 6, 15, and 16 under 35 USC § 103(a) as being obvious in view of U.S. Patent Application Publication 2005/0065753 (Bigus), U.S. Patent Application Publication 2003/0039352 (El-Fakih) and in further view of U. S. Patent Application Publication 2004/0153823 (Ansari). The Appellants ask the Board of Patent Appeals and Interferences (referred to hereinafter as the Board) to review the rejections set forth by the Examiner.

vii. **ARGUMENT**

The Appellants begin these remarks by discussing claim 1. Claim 1 recites a telecommunications system having a control system and peer communication devices. The peer communication devices collect performance data responsive to handling telecommunications data, and transfer the performance data to the control system. The control system processes the performance data from each of the peer communication devices to generate a performance file that indicates the performance of each of the peer communication devices, and transfers the performance file to the peer communication devices. Each of the peer communication devices then processes the performance file to compare its performance to the performance of the other peer communication devices to detect a fault. Responsive to detection of the fault, one (or more) of the peer communication devices processes the performance file to identify a recovery action, and performs the recovery action to attempt to cure the fault.

Thus, the system in claim 1 allows the peer communication devices to *evaluate their own performance* based on the performance file sent by the control system. More particularly, a peer communication device processes the performance file to compare its performance to the performance of the other peer communication devices to detect a fault, and performs a recovery action to attempt to cure the fault. This is what the Appellants refer to as “distributed monitoring”, because the peer communication devices are monitoring their own performance based on the performance file provided by the control system.

The Examiner rejected claim 1 under 35 USC § 103(a) as being obvious in view of U.S. Patent Application Publication 2005/0065753 (Bigus) and U.S. Patent Application Publication 2003/0039352 (El-Fakih). The Appellants disagree, and will show that the cited references fail to teach multiple limitations of claim 1.

Before looking at the individual limitations, the Appellants would like to give the Board a high-level view of claim 1 versus the cited references. As discussed above, the peer communication devices send performance data to the control system, and the control system generates a performance file that indicates the performance of each of the peer communication devices. The control system then sends the performance file to the peer communication devices, which they use to detect their own internal faults. The Appellants generally illustrate this system in the figure below.




FIG. A

As shown in FIG. A, each peer communication device sends performance data to the control system. The control system then generates a performance file that indicates the performance of each of the peer communication devices, and sends the performance file to the peer communication devices. Each of the peer communication devices then processes the performance file to compare its performance to the performance of the other peer communication devices to detect a fault, which provides for distributed monitoring in the network. Distributed monitoring means that the communication devices themselves are monitoring their own performance (i.e., detecting faults based on the performance file) based on the performance file. This is different than prior systems that used a centralized server to monitor performance of the network. The Appellants will show that the cited art does not teach distributed monitoring as recited in claim 1.

Bigus describes a system that monitors the health of a computer system using a centralized server. The computer system includes a server that connects to clients over a network. See FIGS. 1 and 4. Each client includes a monitoring agent that compiles metric data about the client (e.g., CPU usage), and sends the metric data to the server. See FIG. 4 and ¶ [0061]. The server includes a health monitoring system that stores the metric data over a time period. See FIG. 4 and ¶ [0062]. The health monitoring system then analyzes the data to discern fuzzy ranges of normal performance of the clients and the system as a whole. See FIG. 4 and ¶ [0062]. For example, the fuzzy ranges may be averages of the metrics in the system, such as an average CPU usage. The health monitoring system then stores fuzzy sets representing the fuzzy ranges. See FIG. 4 and ¶ [0066]. If the health monitoring system wants to evaluate the

performance of a client or the system as a whole, then the health monitoring system collects current metric data, and compares the current metric data against the fuzzy sets and fuzzy rules. See FIG. 4 and ¶ [0066]. For example, the health monitoring system is able to determine whether the CPU usage in a client is currently above an average usage. The health monitoring system may then provide the results to a system administrator through a GUI. See FIG. 4 and ¶ [0072].

Thus, Bigus describes a centralized server that analyzes the performance of a network, and provides the analysis to a system administrator. The Appellants generally illustrate this system in the figure below.




FIG. B

As shown in FIG. B, the clients send metric data to the centralized server. The server analyzes the metric data to determine how the network is performing, and provides the results to a system administrator. The Appellants point out that there is no performance file sent back to the clients in Bigus, and the clients are not able to detect their own internal faults based on a performance file. Therefore, Bigus does not teach distributed monitoring as in claim 1. It actually teaches away from distributed monitoring by having the centralized server monitor the performance of the network. The Board may refer to FIG. 4 and the corresponding description for a more detailed description of Bigus.

El-Fekih describes a system similar to Bigus, where a centralized server collects performance data from elements in a network. See ¶ [0006-0007]. The centralized server then analyzes the performance data for the network against performance requirements that were

guaranteed to customers. This is done to verify that customers are receiving the level of quality that they expected. El-Fekih also states that a client (customer) may request a report from the service provider verifying that they are receiving the level of quality that they expected. See ¶ [0010]. The report may also be used by the service provider to repair or reconfigure network resources. See ¶ [0010].

Thus, El-Fekih describes a centralized server that receives performance/quality data from network elements, analyzes the performance of the network, and provides a report to a client indicating whether or not the client is receiving the proper QoS. The Appellants generally illustrate this system in the figure below.




FIG. C

As shown in FIG. C, the network elements send performance/quality data to the service management system. The service management system analyzes the performance data to determine if a client is receiving the proper QoS, and provides the results to the client. The Appellants point out that there is no performance file sent back to the network elements in El-Fekih. The network elements in El-Fekih are the devices sending the performance data to the service management system, but they do not receive a performance file back from the service management system. The network elements are not able to detect their own internal faults based on a performance file. Therefore, El-Fekih does not teach distributed monitoring as in claim 1. It also teaches away from distributed monitoring by having the service management system analyze the performance data. The Board may refer to FIG. 1 and ¶¶[0075-0078] for a more detailed description of El-Fekih.

Based on the above remarks, the Appellants submit to the Board that Bigus and El-Fekih teach different types of monitoring than claim 1. This is further illustrated by looking at the individual claim limitations.

First, the Appellants submit that the cited references fail to teach a control system that “processes the performance data from each of the peer communication devices to generate a performance file that indicates the performance of each of the peer communication devices, and transfers the performance file to each of the peer communication devices” as recited in claim 1. As described above, Bigus describes a centralized server that receives metric data from clients, analyzes the metric data to evaluate the performance of the network, and provides the results of the analysis to a system administrator. El-Fekih describes a centralized system that receives performance data from network elements, analyzes the performance data to determine the overall performance of the network, and provides a report to a client (e.g., customer or a service provider) indicating whether or not the agreed-to service levels are being provided. In both Bigus and El-Fekih, performance data for devices are reported to a centralized server, and the centralized server analyzes the performance data to evaluate the performance of the network. The results of the evaluation may then be provided to a system administrator (Bigus), or to a service provider or customer (El-Fekih). However, neither reference describes that a centralized server aggregates the performance data from multiple devices into a performance file, and sends a performance file back to the devices that submitted the performance data in the first place. In Bigus, the client devices send the metric data to the centralized server, but the centralized server in Bigus fails to return a performance file back to these client devices (that submitted the metric data). See ¶ [0061-0066]. In El-Fekih, network elements (34a-f) send performance data to a centralized server (service management system 24), but the centralized server fails to return a performance file back to these network elements (that submitted the performance data). See FIG. 1 and ¶ [0075].

The control system of claim 1 “transfers the performance file to each of the peer communication devices”, where the performance file “indicates the performance of each of the peer communication devices”. Both Bigus and El-Fekih describe communication devices that submit performance data to a centralized server, but neither of these references describes the centralized server transferring a performance file back to each of the communication devices, (where the performance file indicates the performance of each of the communication devices).

Therefore, the combination of these references cannot teach a control system as recited in claim 1.

Secondly, the Appellants submit that the cited references fail to teach that “each of the peer communication devices...processes the performance file to compare its performance to the performance of the other peer communication devices to detect a fault” as recited in claim 1. As described above, neither of the references teaches that a centralized server sends a performance file back to communication devices which submitted performance data to the server. Therefore, the cited references cannot teach communication devices that process a performance file (indicating the performance of each of the communication devices) to detect a fault.

In rejecting this limitation, the Examiner cites to El-Fekih. See page 6 of the final OA. The Examiner’s reliance on El-Fekih is flawed because El-Fekih does not describe a centralized server that returns a performance file to communication devices. Thus, the communication devices in El-Fekih do not have a performance file with which to “compare its performance to the performance of the other peer communication devices to detect a fault”. The Examiner cites to paragraph [0113], which actually teaches away from this claim limitation. Paragraph [0113] in El-Fekih states that:

a service management system may be used to retrieve quality of service information from a network, analyze that information, and compare the analyzed information against defined service or conformance thresholds to determine whether the network is performing up to expectations. If the network is deficient in some way, then a client, such as a service provider or customer, may be notified to allow the client to take corrective action by, for example, reshaping the traffic on the network.

As a reminder, the service management system is the centralized server (see FIG. 1) in El-Fekih which receives performance data from the network elements (e.g., 34a-e), and analyzes the performance data. However, the centralized server does not send a performance file back to the network elements in El-Fekih as recited in claim 1. Because there is no performance file, the network elements in El-Fekih cannot compare their performance against the performance of other network elements to detect a fault. Thus, the Appellants submit that the Examiner failed to show how the cited art teaches this limitation.

Third, the Appellants submit that the cited references fail to teach “at least one of the peer communication devices processes the performance file to identify at least one recovery action, and performs the at least one recovery action to attempt to cure the fault” as recited in claim 1.

In this limitation, not only does a peer communication device evaluate its own performance (in relation to other peer communication devices), but it also performs the appropriate recovery action to cure a fault. The peer communication device does not need to rely on a centralized server (or system administrator) to cure a fault, as it initiates its own recover actions. As described above, neither of the references teaches that a centralized server sends a performance file back to communication devices which submitted performance data to the server. Therefore, the cited references cannot teach communication devices that process a performance file to “identify at least one recovery action” and to “attempt to cure the fault”.

The Examiner has cited to El-Fekih in rejecting this limitation. See page 4 of the final OA. More particularly, the Examiner has cited to paragraph [0113] in rejecting this limitation, but the Appellants submit that the Examiner has misconstrued this paragraph. El-Fekih states that a system may analyze the network performance to determine whether the network is performing up to expectations. See ¶ [0113]. If the network performance is deficient, then a service provider or customer may be notified (such as by a report discussed above). The portion misconstrued by the Examiner is that El-Fekih states that a “client” receiving the report is able to take corrective action by reshaping traffic on the network. The “client” in El-Fekih did not submit performance data to the centralized server. Therefore, the “client” in El-Fekih is not equivalent to a “peer communication device” as recited in claim 1. It is irrelevant whether the centralized server in El-Fekih sends a report to a service provider/customer, and that they reshape traffic on the network. In order to teach claim 1, the centralized server in El-Fekih would have to send a performance file back to the devices which submitted performance data, and these devices would have to detect a fault and perform a recovery action using the performance file. El-Fekih simply does not teach system that sends a performance file to peer communication devices, and that the devices perform recovery actions based on the performance file. Thus, the Appellants submit that the Examiner failed to show how the cited art teaches this limitation.

The Appellants have shown that the combination of Bigus and El-Fekih does not teach the system recited in claim 1. The Appellants submit to the Board that claim 1 is non-obvious over the cited art for at least the reasons set forth above. Similar arguments apply to the other claims.

viii. CLAIMS APPENDIX

1. (Previously Presented) A telecommunication system configured to provide distributed system monitoring, the telecommunication system comprising:

a control system; and

a plurality of peer communication devices, where each peer communication device, responsive to handling telecommunications data, collects performance data and transfers the performance data to the control system;

the control system, responsive to receipt of the performance data from the peer communication devices, processes the performance data from each of the peer communication devices to generate a performance file that indicates the performance of each of the peer communication devices, and transfers the performance file to each of the peer communication devices;

each of the peer communication devices, responsive to receipt of the performance file, processes the performance file to compare its performance to the performance of the other peer communication devices to detect a fault; and

responsive to detection of the fault, at least one of the peer communication devices processes the performance file to identify at least one recovery action, and performs the at least one recovery action to attempt to cure the fault.

2. (Cancelled)

3. (Cancelled)

4. (Cancelled)

5. (Previously Presented) The telecommunications system of claim 1 wherein the at least one peer communication device determines if the fault is cured by the at least one recovery action, generates a report of the fault if the fault is not cured by the at least one recovery action, and transfers the report of the fault to the control system.

6. (Previously Presented) The telecommunications system of claim 5 wherein the control system, responsive to receipt of the report of the fault, identifies at least one recovery action, and performs the at least recovery action on the at least one peer communication device.

7. (Cancelled)

8. (Previously Presented) The telecommunications system of claim 1, wherein:
each of the peer communication devices periodically transfers the performance data to the control system.

9. (Previously Presented) The telecommunications system of claim 1 wherein the performance data includes a performance grade for each of the peer communication devices.

10. (Previously Presented) The telecommunications system of claim 1 wherein the performance file includes a list of performance data for each of the peer communication devices.

11. (Previously Presented) A method of operating a telecommunication system to provide distributed system monitoring, wherein the telecommunication system comprises a plurality of peer communication devices coupled to a control system, the method comprising the steps of:

collecting performance data in each of the peer communication devices responsive to the peer communication devices handling telecommunications data,

transferring the performance data from each of the peer communication devices to the control system,

processing, in the control system, the performance data from each of the peer communication devices to generate a performance file that indicates the performance of each of the peer communication devices,

transferring the performance file from the control system to each of the peer communication devices,

processing the performance file in each of the peer communication devices to compare its performance to the performance of the other peer communication devices to detect a fault; and

responsive to detection of the fault, processing the performance file to identify at least one recovery action, and performing the at least one recovery action to attempt to cure the fault.

12. (Cancelled)

13. (Cancelled)

14. (Cancelled)

15. (Previously Presented) The method of claim 11 further comprising the steps of:

determining if the fault is cured by the at least one recovery action,

generating a report of the fault if the fault is not cured by the at least one recovery action, and

transferring the report of the fault to the control system.

16. (Previously Presented) The method of claim 15 further comprising the steps of:
responsive to receipt of the report of the fault in the control system, identifying at least one recovery action, and performing the at least one recovery action on one of the peer communication devices.
17. (Cancelled)
18. (Previously Presented) The method of claim 11 wherein the step of transferring the performance data from each of the peer communication devices to the control system comprises the step of:
periodically transferring the performance data from each of the peer communication devices to the control system.
19. (Previously Presented) The method of claim 11 wherein the performance data includes a performance grade for each of the peer communication devices.
20. (Previously Presented) The method of claim 11 wherein the performance file includes a list of performance data for each of the peer communication devices.

xi. EVIDENCE APPENDIX

None.

x. RELATED PROCEEDINGS APPENDIX

None.

SUMMARY

The Appellants argue that the Examiner's rejections of claims 1, 5-6, 8-11, 15-16, and 18-20 under 35 U.S.C. 103(a) are inadequate as a matter of law and should be reversed.

Respectfully submitted,

Date: 5-9-2011

/BRETT BORNSEN/

SIGNATURE OF PRACTITIONER

Brett L. Bornsen, Reg. No. 46,566

Duft Bornsen & Fishman, LLP

Telephone: (303) 786-7687

Facsimile: (303) 786-7691

Customer #: 50525